

Evaluating Situation Awareness of Autonomous Systems

Jan D. Gehrke
Center for Computing Technologies (TZI)
Universität Bremen
Am Fallturm 1, 28359 Bremen, Germany
jgehrke@tzi.de

ABSTRACT

Autonomous systems proved to be very successful in specialized problem domains. But their perception, reasoning, planning and behavior capabilities are generally designed to fit special purposes. For instance, a robotic agent perceives its environment in a way that was defined in advance by a human designer. The agent does not exhibit a certain perception behavior because it actually thinks it would be reasonable to do so. But with an increasing level of autonomy as well as a larger temporal and spatial scope of agent operation higher-level situation analysis and assessment become essential. This paper examines criteria for evaluating situation-awareness of autonomous systems and proposes methods to satisfy them. An example application scenario is presented that provides initial results for evaluating situation-aware systems.

1. INTRODUCTION

Autonomous systems are being developed for numerous application areas. These systems proved to be very successful in specialized problem domains, e. g., road driving, area exploration, or robot soccer. Nevertheless, in many cases, the perception, reasoning, planning and behavior capabilities of autonomous systems are *designed* to fit a special purpose. For instance, a robotic agent perceives its environment in a way that was defined in advance by a human designer. Therefore, the agent does not show a certain perception behavior because it actually *thinks* it would be reasonable to do so. It is a reflex rather than a deliberate action. For a lot of applications this might be sufficient. But with an increasing level of autonomy as well as a larger temporal and spatial scope of agent operation higher-level situation analysis and assessment become essential. We will focus on knowledge-based autonomous systems that require up-to-date knowledge to decide on their next course of action. Knowledge-based systems are particularly challenging if they require knowledge of distant locations and/or prediction of future events. These systems cannot rely on their own

sensory capabilities only. They need to infer future states and communicate with other agents to share information.

Situation assessment depends on the current goals of the autonomous system and knowledge about the state of the world. The *relevance* of a specific piece of information or some kind of information for situation assessment is determined by the current plan and other potential plans under consideration. Information relevance leads to the pragmatic dimension of knowledge-based systems, i. e., information on how to use information. Usually, this is implicitly specified by the decision system. For instance, the decision system may use behavior rules whose rule body implies the information needed to evaluate the rule head. Problems arise if the decision-relevant information is not available to the agent. Depending on the applied knowledge representation and reasoning mechanism, missing information is simply considered undefined, default-valued, or false. The latter applies for predicates when using negation as failure [7] as in logic programming. This seems inadequate for autonomous systems in partially observable environments. Even the Open World Assumption turns out to be insufficient for situation assessment. While it prevents to assume wrong states of the unknown world and to base inferences on them, it does not directly enable to reason about what is unknown. As a consequence, the system would not be able to assess a situation correctly, e. g., to detect a harmful risk, because it has a lack of information which it is not aware of.

Hence, autonomous systems doing situation assessment have to be enabled to detect unknown information, thereby becoming *known unknowns*. This detection process must be governed and prioritized by information relevance. If the agent's sensory capabilities cannot provide information needed, other agents or information sources have to be inquired. As an alternative, the agent has to accept its lack of information and address it, e. g., by more cautious behavior.

Only autonomous systems possessing such higher level reasoning capabilities are able to have true *situation awareness* [10]. It even becomes more complicated if agents have to reason about future events. Such agents would need prediction abilities including knowledge on how the world evolves in order to qualify the probability that some information does not change in a given time interval. In spacious environments with mobile robots, additional spatial reasoning about information dynamics is required. This paper outlines a classification of criteria for situation awareness of autonomous systems and suggests knowledge representation and reasoning methods to address them.

The remainder of this paper is structured as follows. Sec-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PerMIS '08, August 19–21, 2008, Gaithersburg, MD, USA
Copyright 2008 ACM 978-1-60558-293-1 ...\$5.00.

tion 2 discusses issues of situation awareness in established autonomous system architectures. Section 3 defines and examines awareness criteria and possible techniques to implement them. A corresponding example scenario for a preliminary study is presented in section 4). The paper concludes with a discussion of this survey (Sect. 5) and a summary (Sect. 6).

2. AUTONOMOUS AGENTS AND SITUATION AWARENESS

Intelligent agents are a fundamental concept in Artificial Intelligence for autonomous decision-making. For most applications domains of agents up-to-date and precise knowledge on the state of the world is crucial for system performance. But surprisingly a lot of usual architectures do not explicitly consider the acquisition of information needed. This limitation might lead to suboptimal, wrong, or even disastrous decisions. Thus situation awareness for autonomous systems intends to evaluate and establish the basis for decision making depending on the agent's current tasks and goals.

Williams et al. [22] have a similar goal. They evaluate the so-called *groundedness* of representations in autonomous systems (mainly those applied in *RoboCup* competitions). The approach defines a measure for the capability of creating and maintaining correct associations between representations and their (physical) real-world counterparts in the system's knowledge base. Proposed qualities of groundedness include, e.g., relevance, accuracy/precision, uncertainty management, and self-awareness w.r.t. the state of the robot body, location, and sensors. The corresponding system evaluation is rather qualitative and performed offline by humans. Thus, the approach provides useful criteria for system evaluation but does not enable the system to reason about itself in order to improve its groundedness.

The simple reflex agent, as the most basic kind of agent, is the worst example for groundedness or situation awareness. Such agents are governed by condition/action rules and always do the same thing given the same perception. In contrast to the model-based reflex agent it has no internal state influencing its decisions [20]. Both reflex agents cannot be considered situation-aware. The simple reflex agent only takes into account instantaneous knowledge; the model-based variant has no notion of information relevance because it has no explicit goals. Anyway, there may be simple but useful tasks that are successfully handled by reflex agents.

The most important term in AI is the rational agent. Wooldridge defines an agent to be "rational if it chooses to perform actions that are in its own best interests, given the beliefs it has about the world" [23, p. 1]. But this definition could also consider an agent rational if it chooses an action without knowing the state of the world. Thus, the situation-aware agent extends and substantiates the classical rational agent definition. The belief about the world is no longer taken for granted but actively controlled by knowledge acquisition as an additional reasoning process.

The belief-desire-intention (BDI) model has become a prevalent approach in academia for deliberative software agent architectures (cf. [23, 15]). It is based on a theory of human practical reasoning developed by Bratman [6]. Human practical reasoning, according to BDI, consists of deliberation, i.e., deciding what state should be achieved, and means-

ends reasoning, i.e., deciding how to achieve it. In the BDI model, an agent is represented by its subjective knowledge about the world (*beliefs*) and persistent goals that should be achieved (*desires*). Desires and current beliefs result in achievable goals and possible actions towards them. Finally, in a process of deliberation, the agent commits to a goal and a corresponding plan (*intention*). The fundamental BDI model does not consider the assessment of beliefs in terms of completeness, correctness/uncertainty, or being up-to-date with respect to the goals to be achieved. Additionally, the model does not take into account different levels of decision making with respect to real-time requirements or temporal scope of action and decision making.

The Real-time Control System (RCS) developed at National Institute of Standards and Technology (NIST) models an intelligent system as a hierarchy of goal-directed sensory-interactive control processes [2] representing organizational levels as well as temporal scopes of decision making. The process hierarchy in RCS enables the decomposition of sub-tasks to different agents as well as different planning intervals within a single agent. Each level contains computational elements for sensory processing, world modeling, value judgment, and behavior generation [1]. Situation awareness could be assigned to the higher-level RCS world modeling components with a tight connection to behavior generation and sensory processing. That is, RCS could be augmented in world modeling by goal-oriented pro-active knowledge acquisition that is governed by behavior generation demands and may provide a focus of attention in sensory processing.

3. CRITERIA FOR SITUATION AWARENESS

Situation awareness is a field of research that commonly examines information requirements of humans for special jobs such as facility monitoring or flying aircraft [10]. Endsley [9, p. 97] describes situation awareness as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future". This leads to Endsley's three levels of situation awareness:

- **Level 1 – Perception:** Basic perception of important information.
- **Level 2 – Comprehension:** Correct interpretation and integration of perceptions as well as relevance assessment.
- **Level 3 – Projection:** The ability to predict future situations based on current perceptions and background knowledge.

Although the definition and the three levels of awareness are intended for human situation awareness they can be adopted for autonomous systems, too. Nevertheless, there are a lot of technical requirements that are partially taken for granted regarding humans but much more challenging for technical systems. Thus, this section proposes the following criteria for situation awareness of autonomous systems and possible methods to fulfill them:

1. **Reasoning about ignorance:** The agent's knowledge base can be queried for unknown or uncertain information.

2. **Model of perception abilities:** The agent is aware of its sensors and the kind of information they may generate.
3. **Model of information relevance:** Based on the current set of goals or a general utility function as well as applied decision rules, the agent can identify information needed and qualify its importance for its performance measure.
4. **Model of information dynamics:** Knowledge on information dynamics that provides stability information for prediction of future states of the world.
5. **Spatio-temporal reasoning:** The model of information dynamics is applied for temporal reasoning on change over time and spatial reasoning on neighborhood and hierarchies of regions
6. **Social ability:** If agents need information beyond their sensory capabilities they have to cooperate with other agents for external information acquisition.

3.1 Reasoning about Ignorance

To enable an autonomous system to measure awareness in a given situation the system needs to know what it does *not* know (*known unknowns*). That is, it has to detect a possible lack of knowledge in its knowledge base. This is particularly important and challenging in environments that are highly dynamic and only partially observable. Unfortunately, many logic-based systems use *negation as failure* in reasoning, i. e., propositions or predicates are assumed to be false if there is no fact or prove stating the opposite. This *closed world assumption* (CWA) is opposed to the *open world assumption* (OWA) which does not make any assumptions about missing knowledge. Instead, logical inference only relies on given facts.

Open world reasoning is particularly applied in ontologies and description logics [3] for concept subsumption in ontology TBoxes, i. e., the schema level of a knowledge base. Situation awareness is rather focused on the instance or assertional level (ABox) of a knowledge base. While the open world assumption also applies for ABox reasoning, it does not provide inferences that would directly provide information on unknown facts. In contrast, a query for the truth of facts that are neither known nor deducible will just return *false*.

As a consequence, CWA and OWA reasoning systems will create a biased view of the state of the world that is very likely to be wrong. If decisions rely on that wrong beliefs system performance is jeopardized. But a strictly logical approach to reason about agent ignorance will raise several difficulties for the logical foundations of representation and reasoning as well as computational efficiency. A structural approach that keeps track of knowledge base changes and instantly replies to queries on ignorance is probably preferable. Anyway, a three-valued logic with OWA reasoning would be required, too.

In general, a strictly logic-based knowledge representation in dynamic environments is debatable. These approaches do not sufficiently handle the uncertainty that is an inevitable consequence in such domains. Though logics are still very useful and powerful for representing background knowledge of an autonomous system, dynamic environmental properties should rather be represented by probabilistic approaches

with aleatory variables. In particular, Bayesian inference provides powerful means to reason about uncertainty. A corresponding approach for value measurement of missing information is presented in Sect. 3.3.

3.2 Model of Perception Abilities

While it is important to be aware of ignorance, agents might deliberately or of necessity choose to decide instantly although having incomplete or uncertain knowledge. This might be due to limited resources or perception abilities, i. e., sensors. An autonomous system that desires a certain kind of information may not be able to acquire this information from its own sensors or data sources. In such cases, it is obviously not reasonable to wait for that information before the decision for which it is considered helpful.

But deliberate ignorance requires background knowledge on what information can be obtained or not. This knowledge can be provided by a sensor model that describes sensors by the type of information they may deliver as well as precision, accuracy, spatial range, and response time. A related ability, the projection of future events, is discussed in Sect. 3.4.

For situation-aware agents we created an OWL-DL [4] ontology of possible sensors and their properties. This ontology does not describe concrete sensor, e. g., some special LIDAR product. Instead, it represents

- types of sensor devices (simple and complex),
- the (usually physical) quantities that are measured,
- the unit of measurement (e. g. SI units),
- the method of measurement (mechanical, electrical, optical, chemical etc.), and
- the sensor quality in terms of measurement errors, precision/accuracy, range, response time etc.

Examples of modeled sensors are those for electrical quantities (i. e. ohmmeter, galvanometer, voltmeter, etc.), electromagnetic radiation sensors for a given spectrum (i. e. infrared, visible light, etc.), acoustic sensors, thermal sensors and range sensors etc. The basic sensors were called *detectors*. But sensors as well as measured quantities can also be more complex. For instance, a weather “sensor” provides more than just one quantity and may aggregate raw measurements to some qualitative information.

The general advantage for agents using this ontology is that they can logically infer whether they have a sensor (or other information source) that will provide a desired information on some quantity or other environmental feature. Additionally, an ontology-grounded value description will allow for better interpretation of sensor data. The ontology also enables agents to communicate about desired information abstracting from the specific sensors used (cf. Sect. 3.6).

3.3 Model of Information Relevance

In large and dynamic environments there will be a lot of information that is inaccessible or unknown to the autonomous system. Other information is available but there are no sufficient resources for processing. Obviously this also holds for humans but they perform very good in a lot of tasks anyway. This is because the major part of the state of the world is simply irrelevant. Humans have the basic

cognitive capability to focus on the relevant stimuli and ignore the others. For low-level tasks this human capability is not based on a deliberation of relevance but basic cognitive functions, e.g., within *superior colliculus* for visual, auditory and somatosensory stimuli. Only relevant stimuli arrive at short-term memory [16].

When trying to imitate this human capability in a cognitive approach to autonomous systems it is practically impossible to design a generic approach for all purposes. Low-level attention control will require low-level implementation because of real-time constraints. But there are also other kinds of information that allow for higher-level assessment of information relevance for determining a focus of attention. Such information is related to decisions with temporal scopes beyond the servo level with 0.05 seconds plans (cf. [2]).

Based on the current set of goals or a general utility function as well as applied decision rules, the agent could identify information needed and qualify its importance for its performance measure. For goal-oriented systems with rule-based decisions, information (as logical predicate) is relevant if it is contained in the body of some rule whose head is going to be evaluated. While this could be identified in a quite straightforward way on syntax level, such a system would require to determine predicates with unknown truth value. This again, has to apply some reasoning on ignorance (Sect. 3.1).

Utility-based systems with probabilistic knowledge representation need to assess information relevance differently. Here we consider a decision-theoretic agent that tries to maximize its utility based on expected utility of each action A with possible outcomes O_1 to O_n of different probabilities given a set of evidences E . In simple cases there is a single decision (not a sequence) and a corresponding action with a discrete distribution of results. Then, the expected utility of action A is defined by:

$$EU(A|E) = \sum_{i=1}^n P(O_i(A)|Do(A), E) \cdot U(O_i(A))$$

Here the best action $\alpha = \max_A EU(A|E)$ depends on the given evidence E . Thus α will potentially change given additional evidence E_j :

$$\alpha_{E_j} = \max_A EU(A|E, E_j)$$

Intuitively, the relevance (or value) V_{E_j} of a new evidence E_j is then defined as the difference in expected utility of the chosen action after and before knowing E_j . But because the agent needs to evaluate the value before acquiring the information it will need to average over possible assignments e_{jk} of E_j , i. e.,

$$V_{E_j} = \left(\sum_k P(E_j = e_{jk}|E) \cdot EU(\alpha_{e_{jk}}|E, E_j = e_{jk}) \right) - EU(\alpha|E)$$

This approach is proposed in *information value theory* by Howard [13]. If the acquisition of information E_j is costly it is only worthwhile if its price is less than its value (presupposing utility is measured in same quantity as price). Note that the information value formula presupposes perfect information. Uncertain information can be included by modeling a probabilistic relation from the uncertain variable to the actual variable.

In order to apply this relevance measurement to assess and potentially increase situation awareness an agent will




Figure 1: Simple decision network example for a route selection problem (adapted from [20, p. 598])

need to proceed in a systematic way. That is, there is no sense in evaluating relevance of all possible information (i. e., aleatory variables). The agent needs some background knowledge on utility-influencing factors.

This knowledge can be provided as a simple list or, more elaborately, as a decision network (or decision diagram) [14]. A decision network combines Bayesian inference with actions and utilities as additional types of nodes besides the usual chance nodes. The factors directly influencing agent utility are represented as parent chance nodes. Possible outcomes of actions are modeled as child chance nodes of each action respectively. There may be additional chance nodes that influence other chance nodes as in Bayesian networks. An example is given in Figure 1.

This decision problem representation enables a systematic ordering of information evaluation and acquisition by traversing the network from the utility node(s) along the path with highest influence on utility. For this purpose, it is necessary to distinguish sensor nodes as special chance nodes since only evidence from sensor nodes is actually available.

Unfortunately, the design of decision networks features the same difficulties as in Bayesian networks: one needs to assign prior probabilities as well as conditional probabilities based on sufficient experience. Because this is not given in general, the agent should be able to update the Bayesian influence model based on ongoing experience. Updates should primarily concern the assignment of probabilities. Creating new Bayesian links between chance nodes is likely to fail due to insufficient statistical evidence in complex and dynamic environments causing possible post-hoc or cum-hoc fallacies.

3.4 Model of Information Dynamics

Whereas decision-theoretic information-gathering agents can be implemented based on decision networks there are special requirements in dynamic environments with decisions of long temporal and spatial scope. As an example, we will use a simple route planning task with a truck that aims at finding the fastest route between two locations in a road network. A decision network for that problem could include chance nodes for influencing factors such as construction work, weather, traffic, time of day, and others (cf. Fig. 1).

The special challenge is that these factors describe future states of the world as well as properties of distant locations.

Consequently, the relevant information is not available from own sensors (distant information) or even theoretically unavailable (future states). Anyway, background knowledge on information dynamics can help infer or estimate such information. If the agent knows that some environment property (e. g., weather) holds in an area of certain extension (e. g. a particular stretch of a freeway) it is easily able to infer from a property of one location to the same property at another location within that area (cf. Sect. 3.5). It gets much more complicated when the agent has to infer this property in an neighboring area. There will be no general probabilistic relation here. It may depend on the property, location, time, and possibly several other factors.

On the other hand, an environment property that is stated to hold for a larger area should allow to deduce this property for sub-areas based on known region paronomies. Again, this inference will provide uncertain information. The statement for a large area will usually only average all locations therein. Thus, single locations might be very different.

Similar considerations also hold for the temporal dimension. A situation-aware system needs prediction abilities to estimate probabilities of variable assignments at future states or it has to acquire reliable predictions from other sources. Obviously, this becomes less achievable when the predicted state is far in the future. On the other hand, there are variables that will not change or only change slowly or rarely.

Re-iterating the above examples, weather or traffic feature some expected stability and transition probability from one state to another, respectively. Such dynamics are often modeled with Hidden Markov Models (HMMs), Kalman filters, or, more general, dynamic Bayesian networks (DBN). That is, one defines an influence or sensor model and a time transition model for the property. Classical DBNs presuppose a fixed discretization of time from one state to another. But in general a reasonable time discretization for a Markov transition model will depend on an average stability of the modeled property. Additionally, information stability may again depend on other factors and change from state to the following.

While DBNs try to infer from sensor evidence and a variable assignment at one state to that of a following (or previous) state, situation-aware agents might not need such kind of prediction for rather distant decisions (i. e. several hours). In some cases these predictions, if possible, will tend to be very uncertain anyway. The most naïve autonomous system will just assume that everything that holds presently will stay that way all along. A situation-aware system should at least consider that the world is dynamic and have some idea about these dynamics. It might be sufficient to know the expected stability (or half-life-period) of each property/variable in order to evaluate the probability that the current value will persist until the desired time in the future. If such persistence is evaluated to be unlikely (i. e. below a predefined threshold) the system is forced to rely on prior probabilities.

In general we advocate the inclusion of dynamic Bayesian networks in decision networks, i. e., dynamic decision networks.

3.5 Spatio-temporal Reasoning

If the situation-aware autonomous system evaluates relevance of missing information it will also need to qualify time




Figure 2: Spatio-temporal region defined by vehicle motion on a road stretch with uncertain driving time.

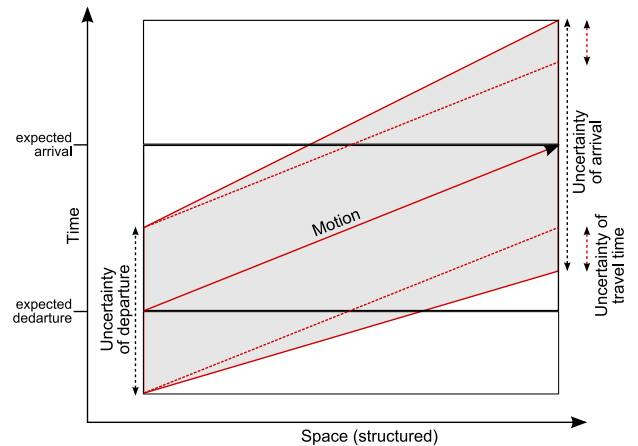


Figure 3: Spatio-temporal region with uncertain departure time.

and space when/where the desired property value holds. In the truck agent routing example the weather or traffic at some distant freeway stretch will influence the expected driving time. When trying to acquire the information from some external source or sensor for that region, the agent also has to consider the expected time when it plans to arrive there. Unfortunately, the arrival gets more and more uncertain with increasing distance of that location.

Figure 2 depicts the situation for a road stretch at the beginning of a truck tour. While the departure time is known arrival time is uncertain. The uncertainty of arrival time or driving time can be provided by a decision network (or regular Bayesian network). But the agent needs to determine the expected spatio-temporal region it will move in to qualify the information for that freeway in space in time. We will call this a *region of relevance* for that agent.

We propose to choose a threshold as a quantile in probability distribution of arrival time to specify a confidence interval of arrival time. The upper bound of that interval will determine the upper temporal bound of the region of relevance for that freeway stretch. All spatio-temporal information overlapping that region is considered potentially relevant w. r. t. value of information (Sect. 3.3).

These regions of relevance will expand in temporal directions for distant locations because uncertainty of departure time increases with uncertainty of arrival time for previous regions. A region of relevance for distant locations is shown in Figure 3.

For qualitative topological inference on this regions the agent can apply usual spatial calculi such as *Region Connection Calculus* (RCC) [19]. This is of particular interest when agents want to share information and need to determine whether there is potential for cooperation. The regions of relevance are then complemented by regions of expertise, i. e., regions other agents can provide information for.

3.6 Social Ability

Social ability is one of the minimal criteria for intelligent agents defined by Wooldridge and Jennings [24]. In this context, social ability refers to communication with other agents. For situation-awareness communication is needed when there is relevant information that is neither available from own sensors nor inferable by reasoning. For instance, information on distant locations needs to be acquired from external sources.

In practice, the mere ability of communication will not suffice. A situation-aware system must determine that some information cannot be acquired from own sensors (Sect. 3.2) but is available from other sources instead. This includes the discovery of such sources, a common communication language with agreed content semantics, and price negotiations if information is subject to charges.

4. APPLICATION SCENARIO

One particular application domain in our research has been *autonomous logistics* [21]. That means, there are agents as autonomous decision makers that represent one or more logistic objects (e. g. truck, container, or pallet) and control logistic processes locally. This usually requires cooperation with other logistic agents as well as continuous acquisition of information for decision-making. This information is provided by local sensors (e. g., humidity sensors for monitoring of perishable goods) as well as external data sources (e. g. for weather and traffic reports).

When developing an autonomous logistics system, we evaluate agent behavior by stochastic simulation before deployment. That is, the developed agents as well as test mocks representing external agents and other entities run in a simulated logistics environment before they are released to a real-world environment. For this purpose we implemented the multiagent-based simulation system PlaSMA (*Platform for Simulations with Multiple Agents*). PlaSMA is based on the FIPA-compatible multiagent development environment JADE [5]. Multiagent-based simulation (MABS) applies object-based modeling [18, 8] as well as discrete time, distributed simulation techniques [11]. In MABS, the agents are parallel logical simulation processes that, in general, represent objects of the simulation model at the same time. Thus, MABS is a natural way to test agent behavior and interaction.

We applied this simulation platform for studies of agent situation awareness in vehicle route planning (similar to the example in section 3). Uncertainty in vehicle route planning is an everyday problem in transportation over longer distances. The shortest route is not always the fastest. Even when considering maximum allowed or average expected speed on single roads the planned routes may prove to be suboptimal. While suboptimal solutions are a natural property in dynamic, open environments with partial observability, usual route planning does not make use of much up-to-date or background information that would be obtain-

able and correlates with travel time.

We conducted simulation studies [12] that help find the utility or relevance of environmental information on travel time and its application in route planning cost functions. Furthermore, the experiments should provide evidence for the time in the future when incoming information is useful and the robustness when working with uncertain or wrong information.

Inductive machine learning was applied to pre-processed traffic information in order to learn traffic models for specific roads depending on weather, time of day, and day of week [12]. The data was based on German and Austrian traffic censuses. The learned traffic model includes prediction rules like

```
[Speed=60]
  <= [Day=Mo] [Time=morning]
     [Weather=moderate..good]
```

That is, average possible speed on the corresponding road is expected to be 60kmph on Monday mornings if weather conditions are from moderate to good. Such predictions were integrated in a truck agent's utility function for best routes in a transportation process. As part of the agents knowledge acquisition component, the planning system initiates the gathering of weather information within a certain distance towards its destination. This information from external weather services together with time and date information is used to consult the rule base for a travel time prediction. The time and date in these rule base queries again depends on prediction of travel time to reach the location of interest (cf. discussion in Sect. 3.5).

In several stochastic simulations these predictions turned out to be valuable. Situation-aware vehicle agents were up to 6.3 % faster on average than regular, non-predictive agents. Additionally, the standard deviation of travel time was reduced by 28 %. With significance level $\alpha = 0.05$ the simulated results to not differ more than 0.0016 % from the actual values. Thus, dedicated means to increase situation-awareness are actually shown to be of advantage in the presented logistics scenario.

5. DISCUSSION

The survey on situation awareness in this paper identifies criteria for knowledge representation and reasoning of autonomous systems in dynamic environments. In particular, the paper focuses on domains with decisions that may have a larger spatio-temporal scope. We do not claim that the means proposed for implementing these criteria are imperative. Furthermore, a reference implementation of the described approach is work in progress. The presented application example still relies on some pre-designed information acquisition behaviors instead of fully goal-oriented and deliberate acquisition. An extended conceptual architecture that includes explicit information acquisition for situation assessment was proposed in [17].

The general question arises whether situation awareness turns out to be necessary or even futile for particular autonomous systems. Furthermore, most systems will probably only fulfill a subset of all criteria. Some criteria are probably not imperative for some degree of situation awareness. Endsley [10] also distinguishes three levels for human situation awareness that could match a particular subset

of the criteria proposed here. Relevance assessment is certainly mandatory for level 2 (*comprehension*). Representing and reasoning with information dynamics pertains to level 3 (*projection*). But social ability (i.e. information exchange) is beyond the scope of Endsley's approach. This criteria could establish a fourth level of awareness. On the other hand, a system could be able to share information but not be capable of projection. So it does not fit in as an extension of level 3.

But assuming an agent that is matching all criteria: Is such an agent capable of actually *measuring* (i.e. quantifying) its situation awareness? We think that information value theory has the answer. The difference between the expected utility with current sensory evidence and the expected utility given complete sensory evidence (considering perception abilities) could provide a corresponding measure. But future research has to be done for a complete theory of situation awareness and its assessment.

6. CONCLUSION

This paper provides a survey on situation awareness for autonomous systems by analyzing features and limitations of existing architectures and proposing a set of criteria to be satisfied by situation-aware agents. The proposed criteria build upon the definition of human situation awareness by Endsley [9, 10] but take into account the special requirements and needed capabilities of technical systems. The proposed criteria are *reasoning about ignorance, an agent perception model, information relevance assessment, a model of information dynamics* and associated *spatio-temporal reasoning capabilities*, and *social ability* for acquiring information beyond own sensors.

The presented logistics routing scenario shows that seemingly barely relevant information on environmental properties can significantly increase performance of autonomous agents. Future work will address a theoretical foundation of situation awareness and its measurement based on dynamic decision networks and information value theory.

7. ACKNOWLEDGMENTS

This research is funded by the German Research Foundation (DFG) within the Collaborative Research Center 637 "Autonomous Cooperating Logistic Processes: A Paradigm Shift and its Limitations" (SFB 637) at the University of Bremen, Germany.

8. REFERENCES

- [1] J. Albus. A reference model architecture for intelligent unmanned ground vehicles. In *Proceedings of the SPIE 16th Annual International Symposium on Aerospace / Defense Sensing, Simulation and Controls*, pages 303–310, Apr. 1–5 2002.
- [2] J. Albus, T. Barbera, and C. Schlenoff. RCS: An intelligent agent architecture. In R. M. Jones, editor, *Intelligent Agent Architectures: Combining the Strengths of Software Engineering and Cognitive Systems*, number WS-04-07 in AAAI Workshop Reports. AAAI Press, 2004.
- [3] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook. Theory, Implementation and Applications*. Cambridge University Press, 2nd edition, 2007.
- [4] S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, L. A. Stein, and F. W. Olin. OWL web ontology language reference. Available from <http://www.w3.org/TR/owl-ref/>, Feb. 10 2004.
- [5] F. Belfemine, A. Poggi, and G. Rimassa. Developing multi-agent systems with a FIPA-compliant agent framework. *Software-Practice and Experience*, 31(2):103–128, Feb. 2001.
- [6] M. E. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, USA, 1987.
- [7] K. L. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 293–322, New York, 1978. Plenum Press.
- [8] P. Davidsson. Multi agent based simulation: Beyond social simulation. In *Multi-Agent Based Simulation, Second International Workshop (MABS 2000)*, number 1979 in LNCS, pages 97–107, Boston, MA, USA, July 8–9 2000. Springer-Verlag.
- [9] M. R. Endsley. Design and evaluation of situation awareness enhancement. In *Proceedings of the Human Factors Society 32nd Annual Meeting*, volume 1, pages 97–101, Santa Monica, CA, USA, 1988. Human Factors Society.
- [10] M. R. Endsley. Theoretical underpinnings of situation awareness. A critical review. In M. R. Endsley and D. J. Garland, editors, *Situation Awareness Analysis and Measurement*, pages 3–32. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 2000.
- [11] R. Fujimoto. *Parallel and Distributed Simulation Systems*. John Wiley & Sons, New York, NY, 2000.
- [12] J. D. Gehrke and J. Wojtusiak. Traffic Prediction for Agent Route Planning. In *International Conference on Computational Science (ICCS 2008)*, number 5103 in LNCS, pages 692–701, Kraków, Poland, June 23–25 2008. Springer-Verlag.
- [13] R. A. Howard. Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, SSC-2(1):22–26, Aug. 1966.
- [14] R. A. Howard and J. E. Matheson. Influence diagrams. *Decision Analysis*, 2(3):127–143, Sept. 2005.
- [15] S. Kirn, O. Herzog, P. Lockemann, and O. Spaniol, editors. *Multiagent Engineering: Theory and Applications in Enterprises*. Springer-Verlag, 2006.
- [16] E. I. Knudsen. Fundamental components of attention. *Annual Review of Neuroscience*, 30:57–78, July 2007.
- [17] M. Lorenz, J. D. Gehrke, J. Hammer, H. Langer, and I. J. Timm. Knowledge management to support situation-aware risk management in autonomous, self-managing agents. In *Self-Organization and Autonomic Informatics (I)*, volume 135 of *Frontiers in Artificial Intelligence and Applications*, pages 114–128, Amsterdam, 2005. IOS Press.
- [18] H. V. D. Parunak, R. Savit, and R. L. Riolo. Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users' Guide. In *MABS 1998*, pages 10–25, Paris, France, 1998. Springer-Verlag.
- [19] D. A. Randell, Z. Cui, and A. G. Cohn. A spatial logic

- based on regions and connection. In B. Nebel, C. Rich, and W. Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, pages 165–176. Morgan Kaufmann, 1992.
- [20] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ, USA, 2nd edition, 2003.
- [21] B. Scholz-Reiter, K. Windt, and M. Freitag. Autonomous logistic processes - New demands and first approaches. In *Proceedings of the 37th CIRP-International Seminar on Manufacturing Systems*, pages 357–362, Budapest, May 19–21 2004.
- [22] M.-A. Williams, P. Gärdenfors, A. Karol, J. McCarthy, and C. Stantom. A framework for evaluating groundedness of representations in systems: From brains in vats to mobile robots. In *IJCAI-05 Workshop on Agents in Real-Time and Dynamic Environments*, pages 17–24, Edinburg, UK, July 30 2005.
- [23] M. Wooldridge. *Reasoning about Rational Agents*. The MIT Press, 2000.
- [24] M. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.